PE3101P: Decision Theory and Social Choice

Lecture 1 — Introduction

Tomi Francis

13 August 2026

1 Course logistics

1.1 Introductions

I'm your instructor for PE3101P: Decision Theory and Social Choice. Feel free to call me Tomi: it's not necessary to use "Dr. Francis" or "Prof. Francis" or whatever.

1.2 Register

As the University wants me to keep records of your attendance, please sign the available printed register when you come in. If your name isn't already on it, please add it. If you do not sign the register, you will not be recorded as being present for the class in question.

1.3 Contact

The best way to reach me is **Email**: tomi.francis@nus.edu.sg

1.4 Office Hours

Office hours: Thursdays, 4-6pm, AS7-03-12. (In practice, this is directly after class, so there's probably no need to wait until 4 pm if you need something from me.)

Hopefully most matters can be resolved either by email or in office hours, but if you need to meet me in-person outside office hours, just send me an email to arrange it.

Note that my office may be moved before the end of this semester; I'll let you know if this happens.

1.5 What the course covers (very short version)

We're going to cover decision-making in two related kinds of case:

(1) Cases where we're not sure what will be the results of our actions, and need to weigh up the risks

(2) Cases where different people want different things, and we need to weigh up these competing interests and decide on behalf of everybody.

Case (1) is decision theory. Case (2) is social choice theory. We'll spend roughly half the semester on each one, with decision theory first.

1.6 Where the course materials live

The main place you should go for course materials is my website: `https://tomifrancis.com/pe3101p.html`

I will try to update canvas where possible, but your primary source should be my website, and in general you can expect the availability of course materials to be better there.

You can expect new materials on my website each week, usually lecture materials, optional assignments, solutions, etc. In general I will try to update the week's materials by the end of that Friday. (So, you get a lecture on Thursday, and the new materials associated with that lecture hopefully go up by the end of the following day.)

For example, notes for this first lecture (the ones you're reading now!) are available on my website. Please refer to them if you find yourself forgetting anything.

If you have any trouble accessing course materials, please send me an email.

1.7 The syllabus

You have the syllabus: a printed copy should be in front of you. If not, please get one from the front.

The syllabus is also available on my website.

I will make it available on canvas when I can, but the copy on my website is more likely to be up to date.

If you've got a logistical question about the course, please reread the syllabus before sending me an email, because it may well include the answer to your question.

1.8 The textbook

The main textbook we will loosely follow is:

Wolfgang Schwarz, *Belief, Desire, and Rational Choice* (2022).

This is an open-access textbook. A pdf copy is available on my website. Please try to use this version, so that we're all reading exactly the same copy.

1.9 Readings

There will usually be 1-2 short readings per week. Readings have been selected with an eye to: 1) informativeness, and 2) fun. We've got a few real crackers, so don't miss out!

In general, I expect you to complete the readings *before* the class. This way, you will find the lectures easier to follow and more useful.

Of course that doesn't go for this week, I'm not expecting you to have read anything yet!

Readings can be found in the syllabus and on the reading list. Please check my website.

1.10 Asking questions and discussion

Questions during lectures are welcome at any point. Please raise your hand if you have a question.

It's possible I may decline to call on you if I'm in a hurry to complete the material, but if this happens (and it should hopefully be only rare) it does **not** mean that you shouldn't raise your hand for questions again later.

Tutorials will be more discussion-focused, and I will be looking for greater participation compared to in lectures.

1.11 Use of AI

DO NOT use AI in any way when completing assessed components of the course. **This is cheating.**

You can use AI (or not) for any non-assessed components of the course, but you should be sensible and use your judgement to use AI in ways that help you learn, and avoid using it in ways that inhibit learning.

Good ideas:

– Producing alternative exercise sheets based on the optional assignments I hand out.
– Using a chatbot to discuss the readings and improve your own understanding after reading them through yourself.
– Creating lists of question types and key facts for revision purposes.
– Asking for explanations of concepts you find difficult to understand.
– Scanning additional readings you have found independently to decide whether they are relevant to the course material and worth reading yourself.

Bad ideas:

– Using a chatbot to complete an optional assignment instead of doing it yourself.
– Asking for answers before spending at least 10 minutes trying to work them out yourself.
– Reading AI-generated summaries of the readings instead of reading the full things yourself.
– Anything that involves switching off your brain.

In short: please DO use AI to supplement your learning. Please DO NOT use AI to replace your learning.

2 Course Overview

We're going to spend the first half of the semester on decision theory, and then the second half of the semester on social choice theory.

2.1 Decision Theory

We'll begin by looking at some decision tables, like this:

	It rains	**It doesn't rain**
Take the umbrella	Dry, but carrying it	Dry, but carrying it
Leave the umbrella	Soaked	Dry, hands free

We'll learn to identify the states of nature, the acts, and the final outcomes, and then we'll consider what are known as dominance principles for decision theory, and why we've got to be careful when applying them.

Afterwards, we'll look at assigning probabilities to states of nature, what that means, and what are some of the rules and constraints we should obey when assigning probabilities, as well as some of the justifications for those rules and constraints: why do people think that it is irrational to assign probabilities in "bad" ways that don't obey the rules?

We're then going to look at a famous problem in philosophy which you may have heard of: **Newcomb's problem**, which is basically about what to do in strange cases where you can **expect** that things are going to go badly for you if you perform a particular action, even though you **know** that performing that action is going to make things better for you, not worse.

(If that sounds impossible, stick around and you'll find out what I mean!)

We're then going to look at the centrepiece of the section on Decision Theory, which is called **Expected Utility Theory**. This is a mathematical result that guarantees that if you make choices under risk in certain ways, which many people think are required by rationality, then you have to act **as though** you are "maximising expected utility", which basically means you're choosing as though you've assigned numbers to complete states of the world which tell you how good those states are for you, and you're trying to choose in a way that makes the average for those numbers as high as possible.

Finally, we're going to look at **justifications** for those ways of choosing under conditions of risk, the so-called "axioms of expected utility theory", and in particular we're going to examine the argument that if you don't choose in the way those rules say you should choose, then there are situations in which you can make yourself worse off for no reason, even though you can see it coming. These are called "money pump arguments".

2.2 Social Choice Theory

We'll then open with social choice theory: the theory of how to balance the competing preferences of many people. And we'll start with a famous argument originally due to **John C. Harsanyi**, which says that the right way to choose is to add up everyone's "utilities" - those numbers that we assign to the world to tell us how good it is, if Expected Utility Theory is right - and choose whatever has the highest total of those numbers. We'll see how the argument works, and also why we might worry about some of the premises of that argument.

Next, we'll look at some game theory, which is about what happens when the outcome isn't decided by a single person, but by many people who also take into account what the others are going to do.

We'll then move on to a famous paradox due to Amartya Sen. The idea here is that even a minimal commitment to the idea that we should choose outcomes which make people better off conflicts with the idea of liberal individual rights.

We'll then look at voting theory, going over a famous result about the impossibility of a satisfactory voting system due to Kenneth Arrow, and we'll round off by discussing whether voting could be irrational, given that an individual is highly unlikely to make a difference to the overall outcome.

So, without further ado, let's start on decision theory!

3 Fact-relative and Evidence-relative correctness

3.1 The Miners Problem

Ten miners are trapped in a shaft and threatened by rising water. You don't know whether the miners are in shaft A or in shaft B. You have enough sandbags to block one shaft, but not both. If you block the right shaft, all miners will survive; if you block the wrong shaft, all of them will die. If you do nothing, both shafts will fill halfway with water and one miner will die.

	Miners in A	**Miners in B**
Block A	all 10 live	all 10 die
Block B	all 10 die	all 10 live
Block neither	1 dies	1 dies

Decision matrix reproduced from Schwarz, Belief, Desire, and Rational Choice *(2022), §1.2. The case originates with Kolodny and MacFarlane, "Ifs and Oughts" (2010).*

3.2 What should you do?

Should you block shaft A, block shaft B, or neither?

. . .

What if the miners are in shaft A? What should you do in that case?

. . .

What if the miners are in shaft A *and you still don't know that they are in shaft A*?

3.3 Fact and Evidence-relative correctness

In one sense, if the miners *are* in shaft A, then the correct thing to do seems to be to block A, regardless of whether you know it.

In another sense, blocking shaft A is arguably foolish and irresponsible, even if you happen to be lucky.

In other words, there seem to be two distinct senses in which you ought to do things.

In what we call the *fact-relative sense*, what you ought to do is the thing that will in fact be best. You have no idea what that thing *is*, but that's neither here nor there.

In what we call the *evidence-relative sense*, what you ought to do is the thing that appears to be best according to the evidence that you have. In this sense, it matters a lot whether you know what to do.

Decision theory is concerned with what you ought to do in the second sense. That's the topic: how can we make good, or rational, decisions on the basis of the evidence that we have?

4 Acts, states and outcomes

Let's look again at that decision table.

	Miners in A	**Miners in** B
Block A	all 10 live	all 10 die
Block B	all 10 die	all 10 live
Block neither	1 dies	1 dies

The columns, indicating whether the miners are in shaft A or shaft B, are called "states", or "states of the world", or sometimes (my preference) "states of nature".

These states of nature are supposed to indicate the ways the world is which you're uncertain about, which you can't change, and which determine what's going to happen, depending on what choice you make. For example, if you block shaft A, then whether the miners die or not depends on whether the miners are in shaft A or shaft B. Changing the shaft you block doesn't change where the miners are, and once you've decided to block a shaft, what happens next depends entirely on where the miners are.

The rows, indicating what you do, are the "acts". These are the choices you make, which determine what's going to happen once it has been determined which state of nature is the actual one.

The cells, indicating what happens, are called the "outcomes", or sometimes "final outcomes", or sometimes "possible worlds". You can think of an outcome as just a complete description of a possible result of your decision.

Question: is the result of "Block neither" when the miners are in A **the same outcome** as the result of "Block neither" when the miners are in B?

4.1 Dominance Reasoning

Imagine we have a fourth option available:

	Miners in A	**Miners in** B
Block A	all 10 live	all 10 die
Block B	all 10 die	all 10 live
Block neither	1 dies	1 dies
Block both	all 10 live	all 10 live

"Block both" results in a better outcome than "Block A" when the miners are in B, and the same outcome when the miners are in A. It similarly results in a better outcome than "Block B" when the miners are in A, and the same outcome when the miners are in B.

That is, looking at the states of nature, compared to blocking just one shaft, blocking both shafts is *definitely* at least as good, and *possibly* better. In this case, we say that blocking both "weakly statewise-dominates" blocking either A or B alone.

Now compare blocking both shafts to blocking neither. In this case, regardless of whether the miners are in shaft A or shaft B, blocking both results in a better outcome. So we can say that blocking both "strongly statewise-dominates" blocking neither.

It seems like it's always better to choose something that strongly or weakly dominates another thing. Why take the risk (or certainty) of getting a worse outcome?

4.2 A counterexample?

Now consider the following decision situation, regarding whether to study for an exam:

	Will Pass	**Won't Pass**
Study	Pass & No Fun	Fail & No Fun
Don't Study	Pass & Fun	Fail & Fun

Decision matrix from Schwarz, Belief, Desire, and Rational Choice *(2022), Exercise 1.3.*

Question: Does "Don't study" statewise dominate "Study"?

...

Yes, but the lesson is not that statewise dominance is wrong, but that we have set up the states in a bad way.

Remember what we said before: states of nature are supposed to be things that you *can't change.* But studying for the exam can make a difference to whether or not you pass.

It's plausible that we should choose statewise dominant acts, but only if we have first set up the states in the right way.

At a first approximation, the right way is for the states of nature to be **probabilistically independent** of which act is chosen. (We will see what this means later.)

5 Chances and Credences

What we ought to do usually depends not just on what the states and outcomes are, but also on how likely each state of nature is to obtain.

	The world ends tomorrow	**It doesn't**
Go into debt, wild party	A magnificent last night	A great night, then years of debt
Stay home, play Scrabble	A quiet last night	A quiet night, and no debt

If you knew it was very likely that the world would end tomorrow, having the wild party might be sensible. But in the actual world, where the chance of the world ending tomorrow is very low, staying at home wins.

Throughout this course, the kinds of probabilities we're going to be interested in are called "credences", or sometimes "degrees of belief". These are, roughly, how likely *you think* something is to happen, on the balance of the evidence. When I say "probability", I'm generally going to mean "credence" unless we say otherwise.

It won't matter much for us, but credences are different to other kinds of chances, usually called "objective chances".

Example. Jack has flipped a coin. He's looking at it now - it landed heads. Jill didn't see the result of the coin flip. She's trying to guess whether the coin landed heads or tails.

The *chance* of the coin having landed tails is zero: in fact, it landed heads. But the *credence* Jill should have in the coin landing tails, if it is a fair coin, is 50%.

We're interested in credences rather than objective chances because we might not know the objective chances, and we're interested in good decision-making in light of the evidence actually available to us.

6 Utilities and probabilities

6.1 Utilities

One way to decide what to do, which we'll see in much more detail later, is known as *maximising expected utility.*

A *utility* is just a number attached to each possible outcome, roughly representing how good it is by your lights.

An *expectation* of a number like this is just the average value of that number. (If you have taken some statistics, you might have heard numbers like these referred to as *random variables.* The expectation of and the mean of a random variable are the same thing.)

Example. If I get \$100 if a fair coin lands heads, and \$0 if a fair coin lands tails, then the expected amount of money I will get is \$50.

Notice that the "expected amount" of money is \$50, even though I know for sure I won't get \$50!

6.2 What is maximising expected utility?

	Heads ($p = 0.5$)	**Tails** ($p = 0.5$)
Take the bet	\$100	\$0
Decline	\$50	\$50

It is exactly what it sounds like. If we have assigned utilities to all possible outcomes via what is called a *utility function*, then each act is associated with an expected utility, because when we focus in on one row, we just have a collection of possible outcomes (which are given utilities) and associated probabilities (the probabilities of the state of nature in their column).

For example, we could decide that each dollar is worth just as much as the last one, and set the utility of an outcome to be the number of dollars gained. In that case, if we maximise expected utility, we will be indifferent between the two acts.

Alternatively, we could decide that 50 dollars is worth 50, but 100 dollars isn't worth twice as much - maybe it's only worth 80.

In that case, the expected utility of not taking the bet is still 50, but the expected utility of taking the bet is 40. So if we maximise expected utility like *this*, we should prefer not to take the bet.

Notice that both the first way, the risk-neutral way, and the second way, the risk-averse way, are ways of maximising expected utility. What we actually do depends on how we assign the utilities.

Notice also that I haven't yet given you any real arguments for why anyone *should* maximise expected utility. We will see some arguments like that in several weeks' time.

7 Bayesianism: two constraints on credences

We haven't yet talked about which probabilities we should assign to states of nature: that is, what our credences should be. We won't answer that question fully in this course, but we will look at a partial answer to the question, which is known as Bayesianism.

Probabilism : credences must obey the **probability axioms**.

The Principle of Conditionalisation : they must change in response to evidence by **conditionalisation**.

8 The probability axioms

A bit more terminology first: an *event* is just a set of final outcomes. We include "no outcomes" and "all possible outcomes" as legitimate events. We want to assign credences to events—the credence assigned to an outcome is then just the credence of the singleton event containing only that outcome. (If you're not sure what this means, don't worry about it.)

The probability axioms say that:

(i) The probability of the universal event—of *something* happening—is 1.

(ii) Probabilities can't be negative.

(iii) The probability of *either* of two completely separate events happening is the sum of their probabilities. (E.g., the probability of a rolled die landing 1 *or* 2 is $1/6 + 1/6 = 1/3$).

8.1 A more formal version (optional)

We write Σ for the set of all events X for the universal event, and we write $\mathrm{Cr}(\sigma)$ to denote the credence an agent gives to event σ.

The probability axioms are:

(i) $\mathrm{Cr}(X) = 1$

(ii) $\mathrm{Cr}(\sigma) \geq 0$

(ii) If σ_i are all disjoint, then $\mathrm{Cr}(\bigcup \sigma_i) = \mathrm{Cr}(\sigma_0) + \mathrm{Cr}(\sigma_1) + \ldots$

9 Conditional probabilities

Next week! Here's a sneak peek:

Definition (Conditional probability)**.**

$$P(A \mid B) = \frac{P(A \wedge B)}{P(B)} \qquad \text{provided } P(B) > 0.$$

10 Readings for next week

BDRC, chapters 2 and 3. You may also find it helpful to read chapter 1 to brush up on what we covered this week!